

SPECIFICATION, SYNTHESIS, AND VERIFICATION OF SOFTWARE-BASED CONTROL PROTOCOLS FOR FAULT-TOLERANT SPACE SYSTEMS

Ufuk Topcu

**University of Texas at Austin
101 East 27th Street, STE 4308
Austin, TX 78712-1500**

16 Aug 2016

Final Report

APPROVED FOR PUBLIC RELEASE; DISTRIBUTION IS UNLIMITED.



**AIR FORCE RESEARCH LABORATORY
Space Vehicles Directorate
3550 Aberdeen Ave SE
AIR FORCE MATERIEL COMMAND
KIRTLAND AIR FORCE BASE, NM 87117-5776**

DTIC COPY NOTICE AND SIGNATURE PAGE

Using Government drawings, specifications, or other data included in this document for any purpose other than Government procurement does not in any way obligate the U.S. Government. The fact that the Government formulated or supplied the drawings, specifications, or other data does not license the holder or any other person or corporation; or convey any rights or permission to manufacture, use, or sell any patented invention that may relate to them.

This report is the result of contracted fundamental research deemed exempt from public affairs security and policy review in accordance with SAF/AQR memorandum dated 10 Dec 08 and AFRL/CA policy clarification memorandum dated 16 Jan 09. This report is available to the general public, including foreign nationals. Copies may be obtained from the Defense Technical Information Center (DTIC) (<http://www.dtic.mil>).

AFRL-RV-PS-TR-2016-0112 HAS BEEN REVIEWED AND IS APPROVED FOR
PUBLICATION IN ACCORDANCE WITH ASSIGNED DISTRIBUTION STATEMENT.

//SIGNED//
RICHARD S. ERWIN
Program Manager

//SIGNED//
PAUL HAUSGEN, Ph.D.
Technical Advisor, Spacecraft Component Technology

//SIGNED//
JOHN BEAUCHEMIN
Chief Engineer, Spacecraft Technology Division
Space Vehicles Directorate

This report is published in the interest of scientific and technical information exchange, and its publication does not constitute the Government's approval or disapproval of its ideas or findings.

Approved for public release; distribution is unlimited.

REPORT DOCUMENTATION PAGE			Form Approved OMB No. 0704-0188	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing this collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.				
1. REPORT DATE (DD-MM-YYYY) 16-08-2016		2. REPORT TYPE Final Report		3. DATES COVERED (From - To) 21 May 2015 – 16 Aug 2016
4. TITLE AND SUBTITLE Specification, Synthesis, and Verification of Software-based Control Protocols for Fault-Tolerant Space Systems			5a. CONTRACT NUMBER FA9453-15-1-0317	
			5b. GRANT NUMBER	
			5c. PROGRAM ELEMENT NUMBER 62601F	
6. AUTHOR(S) Ufuk Topcu			5d. PROJECT NUMBER 8809	
			5e. TASK NUMBER PPM00020204	
			5f. WORK UNIT NUMBER EF125316	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) University of Texas at Austin 101 East 27 th Street, STE 4308 Austin, TX 78712-1500			8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) Air Force Research Laboratory Space Vehicles Directorate 3550 Aberdeen Ave, SE Kirtland AFB, NM 87117-5776			10. SPONSOR/MONITOR'S ACRONYM(S) AFRL/RVSV	
			11. SPONSOR/MONITOR'S REPORT NUMBER(S) AFRL-RV-PS-TR-2016-0112	
12. DISTRIBUTION / AVAILABILITY STATEMENT Approved for public release; distribution is unlimited.				
13. SUPPLEMENTARY NOTES				
14. ABSTRACT In this one-year project, we focused on the control and learning for systems under stochastic uncertainties. We report two sets of results. The first one is motivated by correct-by-construction synthesis for systems with uncertainty in the state due to partial and/or noisy measurements. We developed a new finite-state abstraction technique for such systems. This particular problem was motivated by planning for autonomous space operations. The second one focuses on control and learning in systems in which there is an embedded data-classifier that imperfectly (characterized as stochastically) generates labels from a finite set. Our main contribution was showing how inference techniques for discrete Markov random fields can be applied to learning and control tasks which depend on the output of a noisy classification process.				
15. SUBJECT TERMS Linear Temporal Logic (LTL); Hierarchical Controller Structure; Algorithmic Advances; Proof of Concept; Attitude and Orbit Control Systems				
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT Unlimited	18. NUMBER OF PAGES 16
a. REPORT Unclassified	b. ABSTRACT Unclassified	c. THIS PAGE Unclassified		
				19a. NAME OF RESPONSIBLE PERSON Richard S. Erwin
				19b. TELEPHONE NUMBER (include area code)

(This page intentionally left blank)

TABLE OF CONTENTS

Section	Page
1.0 SUMMARY.....	1
2.0 INTRODUCTION.....	1
3.0 METHODS, ASSUMPTIONS, AND PROCEDURES.....	1
4.0 RESULTS and DISCUSSIONS.....	2
4.1 Filter-Based Stochastic Abstractions for Constrained Planning with Limited Sensing.....	2
4.2 Learning and Control with a Classifier in the Loop	5
5.0 CONCLUSIONS	6
REFERENCES	7
LIST OF ACRONYMS.....	8

ACKNOWLEDGMENTS

This material is based on research sponsored by Air Force Research Laboratory under agreement number FA9453-15-1-0317. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright notation thereon.

DISCLAIMER

The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of Air Force Research Laboratory or the U.S. Government.

1.0 SUMMARY

In this one-year project, we focused on the control and learning for systems under stochastic uncertainties. We report two sets of results. The first one is motivated by correct-by-construction synthesis for systems with uncertainty in the state due to partial and/or noisy measurements. We developed a new finite-state abstraction technique for such systems. This particular problem was motivated by planning for autonomous space operations. The second one focuses on control and learning in systems in which there is an embedded data-classifier that imperfectly (characterized as stochastically) generates labels from a finite set. Our main contribution was showing how inference techniques for discrete Markov random fields can be applied to learning and control tasks which depend on the output of a noisy classification process.

2.0 INTRODUCTION

The demand for reliable, fault-tolerant reactive space systems to deliver sophisticated missions under environmental factors changing over their operational lifespan is increasing their complexity and development costs. These systems are increasingly governed by software-based protocols, both onboard and on the ground, in addition to the conventional physical constraints and actuation and sensing limitations. The interactions between software-based protocols and the underlying physical components call for new methods and tools for the design of complex space systems and for affordably building assurance in their operation. Furthermore, these challenges are not specific only to space systems but are also overarching across an emerging and broad family of systems, called cyber physical systems, including the next generation air traffic control infrastructure, autonomous robots used in transportation and manufacturing, and medical devices as a partial list of concrete examples.

3.0 METHODS, ASSUMPTIONS, AND PROCEDURES

Despite the relatively unique and emerging complications in the control software for space systems, their design often resorts to excessive post-design simulation and tests for establishing trust in their operation. In order to address the lack of and need for systematic methods and automated tools for the design and verification of such systems, this project utilized formal specification languages for unambiguously expressing the correct operation of the system and to develop techniques for both verification against these specifications and synthesis of control logic with accompanying proofs of correctness. We refer to this latter approach as correct-by-construction synthesis [1]. The approach of first specifying and then synthesizing the control protocols---rather than the traditional approach of first designing and then trying to verify---introduces a number of new capabilities. It diffuses the process of building trust throughout the design cycle (rather than being merely a post-design concern) and helps suppress the resulting complexity and the so-called integration surprises. Second, the ability for subsequent synthesis introduces new degrees of freedom to impose a priori architectural constraints and, in turn, to modularize (i.e., separate the concerns to certain extent) the different functionalities in control software.

We take the view in which the operation of control software on space systems is composed of concurrent transitions between modes and focus on the correctness of mode switching sequences and logics. The expected outcomes include principled means for synthesizing and verifying reconfiguration strategies of the system modes to recover from failures, automatically debugging the mode transitions that lead to failures, systematically identifying the cause of known failures, and potentially generating recommendations for possible fixes.

4.0 RESULTS AND DISCUSSION

We now summarize two problems we studied under this effort. The first one is on stochastic abstractions---to be used in high-level, temporal-logic-constrained planning---for systems in which the separation between estimation and control does not hold. In such systems, the properties of the underlying estimation filter have to be taken into account when constructing finite-state abstractions. The second one is on another type of stochastic system in which the stochasticity mainly comes from the embedded machine learning elements (e.g., classifiers). We use a brain-computer interface as the motivating source of stochasticity while imperfections and delays in operator-autonomy interfaces may be interpreted in a similar formalism.

4.1. Filter-Based Stochastic Abstractions for Constrained Planning with Limited Sensing

Motion control problems can sometimes be solved effectively by designing continuous feedback control laws. More often, the complexity of the desired motion or system dynamics leads to the adoption of a hierarchical scheme. In such a scheme, a planning algorithm determines a suitable high-level plan and a continuous controller is designed to implement this plan [2].

The application of planning algorithms for motion planning over continuous spaces often rely on a discretization of the space into finitely many or at least a countable set of states. The discretization may be based on regular grids, random sampling, or the satisfaction of logical proposition. Furthermore, suitable transitions must be defined between the discrete states such that they properly capture the behavior of the original continuous system. The discretization and transitions together form an abstraction of the continuous system. Various decision-making algorithms can be applied to the abstracted system, to generate a high-level plan. Once a suitable path has been planned in the discretized space, the continuous control is tasked with achieving the transitions between the discrete states.

When the uncertainty in the state due to partial and/or noisy measurements is significant, these hierarchical approaches deal with this uncertainty in different ways. The most common approach is to assume that the estimation and control aspects of the problem can be treated separately. For certain systems, such separation is justified due to the existence of a separation principle guaranteeing that the state and estimate will independently converge to their desired values. When such a principle does not exist, this approach may still be reasonable when the estimation error can be reduced arbitrarily fast through design of estimator parameters (e.g. high-gain observers).

Alternatively, the system may be shown to be stable despite large errors in the estimate of the state. These large errors may lead to poor control performance or prevent the control objective from being achieved during the initial time period. As the estimation error decreases over time, eventually the control performance will improve or the control objective will be achieved.

The separation between estimation and control may not always be relied upon in many control problems. When the inputs are constrained, or when the control goals need to be achieved in finite time, the planning methods may need to explicitly account for the uncertainty in the state due to noisy measurement at the instant of implementing the derived control policy. Alternate plans or control strategies may be implemented depending on the amount of information or uncertainty in the system. Solution methods for partially observable Markov decision processes explicitly account for the uncertainty in the state induced by imperfect measurement at the time of choosing actions. The derived optimal control policy is able to handle the trade-off between needing to improve the quality of information and achieving the control objective. However, when attempting to model the continuous system as a partially observable Markov decision process, it is not clear how the convergence properties of the continuous estimation scheme can be exploited.

We developed, in [3], a method of abstraction of a continuous dynamical system with noisy measurements that enables probabilistic-reachability-based methods to be used to predict the probability that a discrete goal state is reached from any other discrete state in the abstraction. These reachability probabilities are valid when the uncertainty in the state is constant. Since the estimation scheme will reduce the uncertainty over time, the policy that maximizes the probability of successfully transitioning between two states at the initial time may be different from that at a later time. We proposed a planning method that uses the predicted levels of uncertainty at future times in order to select the optimal policy at those future times. The result is a time-varying policy which maximizes the lower bound on the probability with which the final discrete state is reached from the initial state. Due to the design of the abstraction and the planning method, a continuous control exists which achieves this lower bound of probability at least.

Let us motivate this procedure through an example. A critical capability for enabling autonomous space operations is autonomous rendezvous and docking. The mission for a given autonomous spacecraft (called the tug) is to rendezvous with an object (called the target) in a given orbit around the Earth. The mission consists of multiple phases during which the sensing capabilities and problem constraints differ. They are the angles-only rendezvous phase, the range-capable rendezvous phase, the docking phase, and the docked phase (see Figure 1).

The distinction between the first two phases lies in the available sensor measurements. The distance between the tug and target cannot be measured reliably using the on-board vision sensor at large distances, however the line-of-sight can be estimated. This is what is meant by the phrase “angles-only”. Once the tug is close enough, the vision sensor can be used to estimate the range, with some noise. This corresponds to the range (distance) capable phase. The docking phase commences once the tug lies within a certain convex region (called the line-of-sight region) close to the target, with low velocity (see Figure 1). The control goal during the first two phases is to ensure that the tug reaches the line-of-sight region within a certain time frame, so that docking can ensue. In Figure 1, the axes represent the frame of the target. The region outside the circle of radius

pr denotes the angles-only rendezvous phase (Phase 1). The green region denotes the range-capable rendezvous phase (Phase 2). The red region is the line-of-sight region which marks the docking phase (Phase 3). An example trajectory of the tug is also shown. The figure is due to [4].




Figure 1. Autonomous Spacecraft Rendezvous and Docking

This problem presents several challenges to designing a suitable controller. The control goal is more complicated than a regulation or tracking condition. The input is constrained, and the measurements are limited and corrupted by noise. The lack of range sensing during the first phase makes planning and prediction complicated. The system is not observable, and control effort must be used in order to obtain a (noisy) estimate of the true state. However, this control effort may decrease the ability of the tug to reach the line-of-sight region within the specified time. Furthermore, the open-loop dynamics are neutrally stable, and care must be taken so that the current control action does not hinder the likelihood of complicating the mission. From the second phase onwards, an estimator can be used to reduce the uncertainty in the state.

Thus, the rendezvous and docking problem motivates the development of planning methods that account for the uncertainty in the state and the behavior of the estimation scheme in order to achieve the control objective with some level of confidence. We create a finite-state abstraction based on the estimation scheme, which can be used to compute bounds on the probability of reaching the goal discrete state based on the uncertainty in the state (represented by a covariance matrix).

4.2. Learning and Control with a Classifier in the Loop

In most dynamical systems, the output of the system is a vector of real numbers obtained through appropriate sensors. The sensed output is used in estimation and feed- back control techniques for achieving various tasks, such as set-point tracking or trajectory tracking. If the sensor is noisy the noise usually takes the form of an additive numerical term centered around the true value. In some systems, the output is obtained using a classification process, and is a member of a discrete set instead of a number. The noise in classification processes results in a feature or data sample being assigned a label different from the true label that should have been assigned to the feature. This type of error is different from the noise in systems with numeric outputs. We aim to understand how to design controllers and learning schemes for systems which have such noisy classifiers in the feedback loop. Figure 2 shows a schematic of a classifier-in-the-loop system. The controller G chooses an action at which drives the finite transition system FTS. The state s and action a are evaluated by the classifier C to produce the class label b , which depends on the task to be completed.




Figure 2. A Classifier-in-the-Loop System Involving a Finite Transition System FTS And a Classifier C

An example where classifier outputs are used in a control task arises in the shared control of semi-autonomous devices using Brain-Computer Interfaces (BCIs) [5]. Brain-Computer Interface technology will enable increased effectiveness of human-machine interaction in general. Current research in BCIs is largely motivated by human-machine interaction in the context of rehabilitation devices. The human nervous system generates signals which can be recorded using techniques such as Electroencephalography (EEG) or Electromyography (EMG). A significant number of BCI approaches extract signals from the human brain by recording brain activity in the form of EEG signals through the human scalp. Using only EEG has the benefit of requiring the user to only wear one piece of headgear to allow interaction with the device. However, extracting meaningful signals using EEG is challenging. Determination of the user's intention is achieved either by training the user to generate a fixed set of signals, or using machine learning techniques to classify recorded EEG signals.

A common task for such devices is to reach a goal state known by the user but unknown to the device, using only classified brain signals as available feedback [6]. The user can either generate brain signals corresponding to control actions, or perform the role of an expert who evaluates the actions selected by the device. Either way, reaching the goal involves solving a learning and control task where a classifier is present in the loop. Techniques for learning which goal is intended by the user are present in the literature. However, no technique to infer when misclassification occurs is attempted.

In our recent work in this project, we outlined a learning and control task in which the information available for feedback was in the form of class labels which are the output of a binary classifier¹. The class labels were used to determine the unknown goal state selected by a human user who evaluates the actions of the device. The evaluation was recorded using EEG signals, leading to the presence of a classifier in the loop. The classification process could make errors, and therefore we developed a method to detect and correct the classification errors with high probability, leading to improved performance in the goal-learning task.

The main insight is that the correct class labels for each state-action pair are related in a predictable way. These relationships can be used to overcome the errors in assigning individual class labels. Moreover, the process of inferring the correct labels can be done efficiently by exploiting existing algorithms for optimization of sub-modular functions.

Our main contribution was showing how inference techniques for discrete Markov Random Fields (MRFs) [7] can be applied to learning and control tasks which depend on the output of a noisy classification process. The inference step yields an estimate of which evaluations by the user were misclassified. The corrected class labels can then be used to provide a better estimate of the goal state.

5.0 CONCLUSIONS

In this one-year project, we focused on the control and learning for systems under stochastic uncertainties. We reported two sets of results. In the first one, we developed a new finite-state abstraction technique to be used in the correct-by-construction synthesis for systems with uncertainty in the state due to partial and/or noisy measurements. Additionally, we reported on our recent results on control and learning in systems in which there is an embedded data-classifier that imperfectly (characterized as stochastically) generates labels from a finite set.

We formulated a series of novel synthesis and verification problems that feature a number of challenging aspects of software-intensive, autonomous space systems, and we proposed preliminary solutions for these problems. Consistent with our original proposal (which outlined a three-year effort while detailing only its first year), our results in this first year lay the ground for a range of new research tasks that are to be addressed in a future effort.

¹ Paper under review for the 2016 Workshop on Algorithmic Foundations of Robotics: H. Poonawala and U. Topcu, “Learning and control with a classifier in the loop”

REFERENCES

- [1] T. Wongpiromsarn, U. Topcu, and R. M. Murray, "Synthesis of control protocols for autonomous systems," *Unmanned Systems*, **1**, pp. 21—29, 2013.
- [2] H. Kress-Gazit, T. Wongpiromsarn and U. Topcu, "Mitigating the state explosion problem of temporal logic synthesis," *IEEE Robotics and Automation Magazine*, **18** (3), pp. 65 – 74, 2011.
- [3] H. Poonawala and U. Topcu, "Filter-Based Stochastic Abstractions for Constrained Planning with Limited Sensing," in the *Proceedings of the 55th Conference on Decision and Control*, Las Vegas, NV, 2016.
- [4] C. Jameson and R. S. Erwin, "A Spacecraft Benchmark Problem for Hybrid Control and Estimation," in the *Proceedings of the 55th Conference on Decision and Control*, Las Vegas, NV, 2016.
- [5] G. Dornhege, J. del R. Millan, T. Hinterberger, D. J. McFarland, and K. R. Miller, "Towards Brain Computer Interfacing," *MIT Press*, Cambridge, MA, 2007.
- [6] J. del R. Millan, "Brain-controlled devices: the perception-action closed loop," *Proceedings of the International Winter Conference on Brain-Computer Interface*, South Korea, 2016.
- [7] D. Koller and N. Friedman, "Probabilistic Graphical Models: Principles and Techniques, Adaptive Computation and Machine Learning," *The MIT Press*, 1st ed., Cambridge, MA, 2009.

LIST OF ACRONYMS

BCI	Brain-Computer Interfaces
EEG	Electroencephalography
EMG	Electromyography
MRF	Markov Random Fields

DISTRIBUTION LIST

DTIC/OCP

8725 John J. Kingman Rd, Suite 0944

Ft Belvoir, VA 22060-6218

1 cy

AFRL/RVIL

Kirtland AFB, NM 87117-5776

2 cys

Official Record Copy

AFRL/RVSV/Richard S. Erwin

1 cy

(This page intentionally left blank)